

WHAT IS CLAIMED IS:

1. A chip-multiprocessing system with scalable architecture, comprising on a single chip:
a plurality of processor cores;
a two-level cache hierarchy including
5 a pair of instruction and data caches for, and private to, each processor core, the pair being first level caches, and
a second level cache with a relaxed inclusion property, the second-level cache being logically shared by the plurality of processor cores, the second level cache being modular with a plurality of interleaved modules;
10 one or more memory controllers capable of operatively communicating with the two-level cache hierarchy and with an off-chip memory;
a cache coherence protocol;
one or more coherence protocol engines;
an intra-chip switch; and
15 an interconnect subsystem.
2. A chip-multiprocessing system as in claim 1, wherein the scalable architecture is targeted at parallel commercial workloads.
- 20 3. A chip-multiprocessing system as in claim 1, further comprising on a single I/O chip (input output chip):
a processor core similar in structure and function to the plurality of processor cores;
a single-module second-level cache with controller;
an I/O router; and
25 a memory that participates in the cache coherence protocol.
4. A chip-multiprocessing system as in claim 1, wherein the plurality of core processors are each a single-issue, in-order processor configured with a pipelined datapath and hardware support for floating-point operations.

30

5. A chip-multiprocessing system as in claim 1, wherein the plurality of processor cores are each capable of executing an instructions set of the ALPHA™ processing core.

6. A chip-multiprocessing system as in claim 1, wherein the plurality of processor cores are each configured with a branch target buffer, pre-compute logic for branch conditions, and a fully bypassed datapath.

7. A chip-multiprocessing system as in claim 1, wherein each of the plurality of processor cores is capable of separately interfacing with either of the instruction and data caches, and wherein each of the caches is configured for single-cycle latency.

8. A chip-multiprocessing system as in claim 1, wherein the interconnect subsystem includes a network router, a packet switch and input and output queues.

9. A chip-multiprocessing system as in claim 1, wherein the single chip creates a node, and wherein the coherence protocol engines include a home engine and a remote engine which support shared memory across multiple nodes.

10. A chip-multiprocessing system as in claim 1, further comprising:
a system control module that takes care of system initialization and maintenance including configuration, interrupt handling, and performance monitoring.

11. A chip-multiprocessing system as in claim 1, wherein each of the plurality of interleaved modules of the second level cache has its own controller, on-chip tag and data storage, and wherein each module is attached to one of the memory controllers which interfaces to a bank of memory chips.

12. A chip-multiprocessing system as in claim 11, wherein each bank of memory chips includes DRAM (dynamic random access memory) chips.

13. A chip-multiprocessing system as in claim 1, wherein the second level cache is interleaved into eight modules.

14. A chip-multiprocessing system as in claim 1, wherein each of the instruction and data
5 caches is a two-way set-associative, blocking cache with virtual indices and physical tags.

15. A chip-multiprocessing system as in claim 1, wherein each instruction cache is kept coherent by hardware.

10 16. A chip-multiprocessing system as in claim 1, wherein each of the second level cache modules includes an N-way set associative cache and uses a round-robin or least-recently-loaded replacement policy if an invalid block is not available.

15 17. A chip-multiprocessing system as in claim 1, wherein each of the plurality of interleaved modules has its own control logic for maintaining intra-chip coherence and cooperation with the plurality of coherence protocol engines, an interface to its dedicated memory controller, and an intra-chip switch interface for intra-chip communication within the single chip.

18. A chip-multiprocessing system as in claim 1, wherein the pair of instruction and data
20 caches includes a first state field per each cache line present therein the first state field having bits related to the MESI (modified, exclusive, shared, invalid) protocol.

19. A chip-multiprocessing system as in claim 18, wherein the second level cache maintains a duplicate of the first state fields from the first-level pairs of instruction and data caches, the
25 duplicate being maintained in order to avoid the need for a first-level cache lookup for cache lines that map to given addresses of corresponding requested cache lines.

20. A chip-multiprocessing system as in claim 18, wherein the second level cache holds a second state field for each cache line present therein, the second state field having bits related to
30 the MESI protocol, wherein the second level cache maintains a duplicate of the first state fields,

and wherein on every second level cache access the duplicate first state fields and the second state fields are accessed in parallel.

21. A chip-multiprocessing system as in claim 1, wherein the single chip creates a node, and
5 wherein information about sharing of data across nodes is kept in a directory in a memory accessed via the memory controllers.

22. A chip-multiprocessing system as in claim 21, wherein the second level cache includes a controller, and wherein manipulation and interpretation of the directory is done by the protocol
10 engines, although the controller also interprets the directory, but merely for determining whether a cache line is cached remotely to the single chip.

23. A chip-multiprocessing system as in claim 1, wherein the interconnect subsystem includes at least one datapath, and wherein the interconnect subsystem is a crossbar configured
15 with a uni-directional, push-only interface, and is capable of scheduling data transfers according to datapaths availability, pre-allocating datapaths, speculatively asserting a requester's grant signal, and supporting back-to-back transfers without dead-cycles between transfers.

24. A chip-multiprocessing system as in claim 11, wherein the controllers in the plurality of
20 interleaved modules are responsible for enforcing coherence within the single chip.

25. A chip-multiprocessing system as in claim 11, wherein access to any of the one or more memory controllers is controlled by and muted through a corresponding one of controllers in the plurality of interleaved modules.

26. A chip-multiprocessing system as in claim 1, wherein the memory controller includes a memory access controller with high speed interface circuitry and a memory controller engine capable of scheduling second-level cache memory access.

27. A chip-multiprocessing system as in claim 1, wherein the coherence protocol engines are implemented as similarly structured microprogrammable controllers, although each of them has its respective microcode.

5 28. A chip-multiprocessing system as in claim 1, wherein each of the coherence protocol engines is configured with an input stage, a microcode-controlled execution stage and an output stage.

29. A chip-multiprocessing system as in claim 1, wherein at least one of the coherence
10 protocol engines is configured to execute protocol code that includes instructions named Send, Receive, Lsend, Lreceive, Test, Set and Move.

30. A method for scalable chip-multiprocessing, comprising:
providing on a single chip

15 a plurality of processor cores,
a two-level cache hierarchy including
a pair of instruction and data caches for, and private to, each processor
core, the pair being first level caches, and
a second level cache with a relaxed inclusion property, the second-level
20 cache being logically shared by the plurality of processor cores, the second level
cache being modular with a plurality of interleaved modules,
one or more memory controllers capable of operatively communicating with the
two-level cache hierarchy and with an off-chip memory,
a cache coherence protocol,
25 one or more coherence protocol engines,
an intra-chip switch, and
an interconnect subsystem,
wherein the single chip creates a node; and
providing one or more than one of the nodes to create, in a modular scalable fashion, a
30 glueless multiprocessor.

31. A method for scalable chip-multiprocessing as in claim 30, further comprising:
providing on a single I/O chip (input output chip)

a processor core similar in structure and function to the plurality of processor
cores,

5

a single-module second-level cache with controller,

an I/O router, and

a memory that participates in the cache coherence protocol.